120

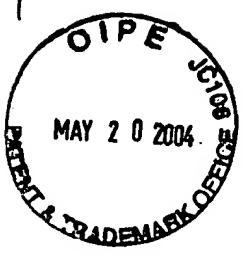

AF/2182

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Atty. Docket No.: **AUS920000447US1**

In re Application of:

ARNDT, ET AL.

Serial No.: **09/766,764**

Examiner: **J. SCHNEIDER**

Filed: **January 23, 2001**

Art Unit: **2182**

For: **DMA WINDOWING IN AN LPAR ENVIRONMENT USING DEVICE ARBITRATION LEVEL TO ALLOW MULTIPLE IOAs PER TERMINAL BRIDGE**

RECEIVED
MAY 2 1 2004
Technology Center 2100

APPEAL BRIEF UNDER 37 C.F.R. 1.192

Mail Stop Appeal Briefs - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Brief is submitted in triplicate in support of the Appeal in the above-identified application. A Notice of Appeal was filed in this case on March 11, 2004 and received in the patent office on March 18, 2004.

CERTIFICATE OF MAILING 37 CFR 1.8(A)

I hereby certify that this Appeal Brief is being deposited with the United States Postal Service via First Class Mail in an envelope addressed to Mail Stop Appeal Brief-Patents, Commissioner of Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450.

Pia R. Lorenzana — May 17, 2004
Date

05/21/2004 AWONDAF1 00000126 090447 09766764
01 FC:1402 330.00 DA

REAL PARTY IN INTEREST

The real party in interest in the present Appeal is International Business Machines Corporation of Armonk, New York, the Assignee of the present application as evidenced by the Assignment set forth at reel 011499, frame 0128.

RELATED APPEALS AND INTERFERENCES

There are no other appeals or interferences known to Appellants, the Appellants' legal representative, or assignee, which directly affect or would be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-6 stand finally rejected as noted in the Examiner's action dated January 9, 2004.

STATUS OF AMENDMENTS

No amendment has been submitted subsequent to the final rejection.

SUMMARY OF THE INVENTION

As set forth in the present specification at page 4, line 3, *et seq.*, input/output adapters utilized by an operating system image in a logically partitioned data processing system are prevented from fetching or corrupting data from a memory location allocated to another operating system image within the data processing system. In one embodiment, the data processing system includes a plurality of logical partitions, a plurality of operating systems, a plurality of memory locations, a plurality of input/output adapters and a hypervisor. Each of the operating system images is assigned to a different one of the logical partitions. Each of the memory locations and each of the input/output adapters is assigned to one of the logical partitions. The hypervisor prevents transmission of data between an input/output adapter in one of the logical partitions and memory locations assigned to other logical partitions during a direct memory access operation by assigning each of the input/output adapters a range of input/output bus DMA addresses. When a request from an OS image to map some of its memory for a DMA operation is received, the hypervisor checks that the memory address range and the input/output adapter are allocated to the requesting OS image and that the input/output bus DMA range is

within that allocated to the input/output adapter. If these checks are passed, the hypervisor performs the requested mapping, otherwise the request is rejected.

The present invention further contemplates the use of terminal bridges to support multiple input/output adapters. In this embodiment, each terminal bridge has a plurality of sets of range registers, each associated with a respective one of the input/output adapters to which it is connected. An arbiter is provided which selects one of the input/output adapters to utilize the PCI bus. The terminal bridge can then examine the grant signals from the arbiter to the input/output adapters to determine which set of range registers is to be utilized.

As illustrated in Figure 3 and as described in the present specification at page 11, line 21, *et seq.*, a block diagram of a data processing system is illustrated. As depicted, the data processing system is implemented as a logically partitioned server, such as server 104 within Figure 1. Data processing system 300 may be a symmetric multi-processor system including a plurality of processors 301, 302, 303, and 304 connected to a system bus 306. Also connected to system bus 306 is memory controllers/cache 308, which provides an interface to a plurality of logical memories 306-363. Input/output bus bridge 310 is connected to a system bus 306 and provides an interface to input/output bus 312. Memory controller/cache 308 and input/output bus bridge 310 may be integrated as depicted.

Data processing system 300 is a logically partitioned data processing system. Thus, data processing 300 may have multiple heterogeneous operating systems running simultaneously. Each of these multiple operating systems may have any number of software programs executing within it. Data processing system 300 is logically partitioned such that different input/output adapters 320-321, 328-329, 336-337 and 346-347 may be assigned to different logical partitions.

Thus, for example, supposed data processing system 300 is divided into 3 logical partitions, P1, P2 and P3. Each of the input/output adapters 320-321, 328-329 and 336-337, each of the processors 301-304, and each of the local memories 360-364 is then assigned to one of the 3 partitions. For example, processor 301, memory 360, input/output adapters 320, 328 and 329 may be assigned to logical partition P1; processors 302-303, memory 361 and input/output adapters 321 and 337 may be assigned to partition P2; and, processor 304, memories 362-363 and input/output adapters 336 and 346-347 may be assigned to logical partition P3.

Each operating system executing within data processing system 300 is assigned to a different logical partition. Thus, each operating system executing within data processing system

300 may access only those input/output units that are within its own logical partition. For example, one instance of the advanced interactive executive (AIX) operating system may be executing within partition P1, a second instance (image) of the AIX operating system may be executing within partition P2 and a Windows2000™ operating system may be operating within logical system P1. Windows2000™ is a product and trademark of Microsoft Corporation of Redmond, Washington.

Peripheral component interconnect (PCI) host bridge 314 is connected to input/output bus 312 and provides an interface to PCI local bus 315. A number of Terminal Bridges 316-317 may be connected to PCI bus 315. Typical PCI bus implementations will support four to ten terminal bridges for providing expansion slots or add-in connectors. Each of Terminal Bridges 316-317 is connected to a PCI input/output adapter 320-321 through a PCI bus 318-319. Each of the Terminal Bridges 316-317 is configured to prevent the propagation of errors up into the PCI host bridge 314 and into higher levels of data processing system 300. By doing so, an error received by any of Terminal Bridges 316-317 is isolated from the shared buses 315 and 312 of the other input/output adapters 321, 328-329, 336-337, and 346-347 that may be in different partitions. Therefore, an error occurring within an input/output device in one partition is not "seen" by the operating system of another partition. In this manner, the integrity of the operating system in one partition is not effected by an error occurring in another logical partition.

Referring now to page 23 of the specification, at line 8 *et seq.*, Figure 7 is described and a detailed description of a terminal bridge is illustrated. As depicted, the control logic of a terminal bridge such as terminal bridge 702 includes an arbiter 714 which controls access to a PCI bus 716. The bus requests signals 710 from the input/output adapters 700 are fed into arbiter 714 which then determines which input/output adapter gets to use the bus and thereafter, arbiter 714 signals that input/output adapter via a grant signal 718. By examining these grant signals 718, terminal bridge 702 can use the appropriate set of range registers 712 that are assigned to that particular input/output adapter. If an input/output adapter receives a grant from the arbiter and the address that the input/output adapter is attempting to use is outside the range indicated by the selected range registers, then the terminal bridge signals the input/output adapter to abort that operation and thus prevents the input/output adapter from accessing memory which that particular input/output adapter is not permitted to access.

In this manner, transmission of data between a particular input/output adapter associated with a particular logical partition is prevented from accessing memory locations assigned to a different logical partition.

ISSUES

I. Is the Examiner's rejection of claims 1-6 under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which the Applicant regards as the invention well founded?

II. Is the Examiner's rejection of claims 1-6 under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent 4,843,541 issued to *Bean, et al.*, in view of U.S. Patent No. 6,438,671 issued to *Doing, et al.* well founded?

III. Is the Examiner's rejection of claims 3, 5 and 6 under 35 U.S.C. § 103(a) as being unpatentable over *Bean, et al.*, in view of *Doing, et al.* as applied to claims 1-2 and further in view of the Admitted Prior Art well founded?

III. Is the Examiner's rejection of claim 4 under 35 U.S.C. § 103(a) as being unpatentable over *Bean, et al.*, in view of *Doing, et al.* and the Admitted Prior Art and further in view of U.S. Patent No, 6,584,530 issued to *Kondo, et al.* well founded?

IV. Is the Examiner's rejection of claim 1 under the judicially created doctrine of obviousness-type double patenting over claim 1 of co-pending Application No. 09/589,665 well founded?

GROUPING OF THE CLAIMS

For purposes of this Appeal, claims 1-6 stand or fall together as a single group.

ARGUMENT

In the Examiner's Action dated January 9, 2004 the Examiner has finally rejected claims 1-6 under 35 U.S.C. §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which the Applicant regards as the invention. That rejection is not well founded and it should be reversed.

The Examiner has based this rejection upon a belief that the term "terminal bridge" as set forth within the claim is "not defined in the specification in such a way as to give clear indication as to what is being claimed." Applicant respectfully disagrees with the Examiner's position and believes that this rejection should be reversed for the following reasons.

As set forth in the present specification at page 13, lines 8 *et seq.*, and as illustrated within Figure 3, multiple terminal bridges are described and depicted in the present specification. Each terminal bridge is described in the present specification as being configured to prevent the propagation of errors up into the PCI host bridge and into higher levels of the data processing. Figure 7 of the present application further illustrates a single terminal bridge 702 and that Figure is described in the present specification at page 23, lines 7 *et seq.* As set forth at the aforementioned portion of the specification, terminal bridge 702 includes an arbiter 714 which controls access to PCI bus 716. Bus request signals 710 from input/output adapters are fed into arbiter 714 which then determines which input/output adapter gets to utilize the bus and thereafter arbiter 714 signals that input/output adapter via a grant signal. The grant signals permit terminal bridge 702 to utilize the appropriate set of range registers 712 which are assigned to that particular input/output adapter. Any attempt by an input/output adapter to utilize a memory address outside the appropriate set of range registers is then denied.

Applicant urges the Board to consider that a "terminal bridge" is clearly described in the present specification, illustrated in great detail within Figure 7 and specified in terms of its functionality. Applicant has also submitted during the prosecution of this application copies of dictionary pages defining the terms "terminal" and "bridge" and shown that the utilization of these two terms in a phrase as illustrated and defined in the present specification is well within the scope of their normal definition and does not do clear violence to the ordinary meaning of those words as recognized by one having ordinary skill in the art.

Applicant urges the Board to consider that a terminal bridge such as terminal bridge 702, as illustrated and described with respect to Figure 7 of the present application, is interposed between PCI host bridge 704 and input/output adapters 700 and receives data on the buses of there between. It is a requirement under 35 U.S.C. §112, second paragraph that the invention be described in a manner such that one having ordinary skill in the art will clearly contemplate the

invention in a manner which particularly points out and distinctly claims the subject matter and Applicant urges that the claims in the present application clearly do so.

The Examiner's position that the phrase "terminal bridge" is entitled to broader interpretation than that set forth within the specification and "could be as little as a wire between two devices" is inconsistent with the proper interpretation of the present claims under the sixth paragraph of 35 U.S.C. §112, as these elements are clearly set forth in the present claims in view of the function performed thereby. Consequently, Applicant urges that the Examiner's rejection of claims 1-6 under 35 U.S.C. §112, second paragraph is improper and reversal of that rejection is believed to be appropriate.

Next, the Examiner has rejected claims 1-6 under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent 4,843,541, issued to *Bean, et al.*, in view of U.S. Patent No. 6,438,671, issued to *Doing, et al.* That rejection is not well founded and it should be reversed.

As set forth within claim 1 of the present application, as an exemplar of the claims of group 1, the logically partitioned data processing system of the present invention is described as a plurality of input/output adapters each one of which is associated with a different one of a plurality of logical partitions which are described. These multiple input/output adapters are connected to a terminal bridge which is interposed between those input/output adapters and the data transmission bus, and which prevents transmission of data between a given one of the multiple input/output adapters and the data transmission bus.

In contrast, *Bean, et al.* describes a partitioned data processing system which utilizes multiple Virtual Machines (VM). As set forth within *Bean, et al.* at col.1, lines 35 *et seq.*, in a Virtual Machine (VM) system, each user is given "an apparent or logical CPU." And "Plural Logical CPUs could share each real CPU resource(s) in the system." (emphasis added). Thus, unlike the claims of the present application, in a Virtual Machines (VM) system all actual resources within the system are accessible by each Virtual Machine (VM). The mirrored utilization of address translation so that portions of the system memory are assigned to a particular Virtual Machine (VM), does not, in the opinion of the Applicant, show or suggest in any way the physical limitation of a plurality of input/output adapters in a manner which is set forth expressly within the claims of the present application, and *Bean, et al.* is believed not to be a valid reference for rejecting the claims of the present application in that *Bean, et al.* fail to

show or suggest in any way the physical limitation of access to one of a plurality of input/output adapters utilizing a terminal bridge, as set forth in the claims of the present application.

The Examiner cites *Doing, et al.,* in view of *Bean, et al.'s* failure to explicitly teach a data transmission bus and a terminal bridge connected to that data transmission bus, noting that the use of buses and bus bridges is well known in the computer art. The Examiner cites *Doing, et al.* and notes that Figure 1 thereof teaches multiple input/output processing units and although *Doing, et al.* does teach the use of multiple input/output processing units nothing within *Doing, et al.* shows or suggests the interposition of a terminal bridge between each of those multiple input/output adapters processing devices and bus interface 105 as expressly set forth within the claims of the present application.

Applicant also notes that *Doing, et al.* issued on August 20, 2002 and that the present application was filed on January 23, 2001. Thus, Applicant believes *Doing, et al.* is a reference against the present application only under the provisions of 35 U.S.C. §102(e)(2) and, *Doing, et al.* is assigned on its face to International Business Machines Corporation of Armonk, New York. As evidenced by the assignment noted above at reel 011499, frame 0128, the present application is also assigned to International Business Machines Corporation and consequently the present invention, at the time the invention was made, was owned by the same person as *Doing, et al.* and consequently, that reference is excluded as prior art under 35 U.S.C. §103(c).

The Examiner has also rejected claims 3, 5 and 6 under 35 U.S.C. § 103(a) as being unpatentable over *Bean, et al.*, in view of *Doing, et al.* as applied to claims 1-2 and further in view of the Admitted Prior Art. This rejection is not well founded and it should be reversed.

As noted above, *Doing, et al.* is not an effective reference against the present application under the provisions of 35 U.S.C. §103(c) and *Bean, et al.* teaches a Virtual Machine (VM) in which all resources are shared by each Virtual Machine (VM). Consequently, the use of DMA range register facilities cannot be said to show or suggest the invention set forth within claims 3, 5 and 6 when combined with *Bean, et al.*, the only effective reference cited by the Examiner. Consequently, this rejection is not well founded and it should be reversed.

The Examiner has also rejected claim 4 under 35 U.S.C. § 103(a) as being unpatentable over *Bean, et al.* and *Doing, et al.* and the Admitted Prior Art, and further in view of U.S. Patent

No, 6,584,530 issued to *Kondo, et al.* That rejection is not well founded and it should be reversed.

As noted above, *Bean, et al.* teaches a Virtual Machine (VM) system in which all actual resources within their system are accessible to each Virtual Machine (VM). Further, *Doing, et al.* is no longer effective as prior art in the present application under 35 U.S.C. §103(c) and *Kondo, et al.* merely teaches the utilization of a bus converter which is coupled to multiple input/output devices. The bus arbiter described as set forth within Figure 1 of *Kondo, et al.* is utilized to prevent low-speed input/output accesses from obstructing transactions with main storage, and does not, in any way, shape or form, prevent transmission of data to locations which are unassigned to particular logical partitions to a particular input/output device as expressly set forth within the claims of the present application and Applicant therefore urges that this rejection is not well-founded and it should be reversed.

The Examiner has also provisionally rejected claim 1 under the judicially created doctrine of obviousness-type double patenting as being unpatentable over claim 1 of co-pending Application No. 09/589,665. That rejection is not well founded and it should be reversed.

The Examiner notes that the co-pending application fails to claim "a data transmission bus and a terminal bridge connected to that data transmission bus..." but goes on to note that such elements are well known in computer systems. Applicant believes that this rejection is contrary to the Examiner's position with respect to the rejection of claims 1-6 under 35 U.S.C. §112, second paragraph, in which the Examiner urges that the term "terminal bridge" is not well defined and that the Examiner may not thereafter choose to interpret that element contrary to the definition and illustrations of the present application as "well known" in the prior art in order to sustain an obviousness-type double patenting rejection over the co-pending application and reversal of the Examiner's provisional rejection is therefore respectfully requested.

CONCLUSION

Please charge the fee of $330.00 for submission of a Brief in Support of Appeal to IBM Corporation Deposit Account No. 09-0447. No additional filing fee is believed to be necessary; however, in the event that any additional fee is required, please charge it to IBM Deposit Account Number 09-0447.

Respectfully submitted,

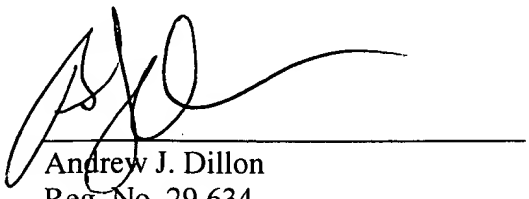

Andrew J. Dillon
Reg. No. 29,634
DILLON & YUDELL LLP
8911 N. Capital of Texas Highway
Suite 2110
Austin, Texas 78759
512-343-6116

ATTORNEY FOR APPLICANTS

APPENDIX

1. A logically partitioned data processing system, comprising:

a plurality of logical partitions;

a plurality of operating systems, each assigned to one of said plurality of logical partitions;

a plurality of memory locations, each location assigned to one of said plurality of logical partitions;

a data transmission bus;

at least one terminal bridge connected to said data transmission bus;

a plurality of input/output adapters, each associated with a different one of said plurality of logical partitions, said input/output adapters being connected to said terminal bridge; and

means for preventing transmission of data between a given one of said input/output adapters which is associated with a first one of the plurality of logical partitions, and memory locations unassigned to said first one of said plurality of logical partitions.

2. The logically partitioned data processing system of Claim 1 wherein said data transmission bus is a PCI bus, and further comprising:

a PCI host bridge connected to said PCI bus; and

an input/output bus connected to said PCI host bridge.

3. The logically partitioned data processing system of Claim 1 wherein said terminal bridge has a plurality of sets of range registers, each associated with a respective one of said input/output adapters.

4. The logically partitioned data processing system of Claim 3 further comprising an arbiter which selects one of said input/output adapters to use said data transmission bus, wherein said transmission preventing means assigns one of said sets of range registers based on a grant signal from said arbiter.

5. The logically partitioned data processing system of Claim 3 wherein said sets of range registers contain direct memory access addresses which limit operations that may be placed onto said data transmission bus by said input/output adapters.

6. The logically partitioned data processing system of Claim 3 wherein said sets of range registers are programmable.